The Quarterly Journal of Mathematics Advance Access published on March 28, 2018
Quart. J. Math. 69 (2018), 1089–1099; doi:10.1093/qmath/hay015

IRREDUCIBLE COMPOSITIONS OF DEGREE TWO POLYNOMIALS OVER FINITE FIELDS HAVE REGULAR STRUCTURE

by ANDREA FERRAGUTI[†]

(*University of Cambridge, DPMMS, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WB, UK*)

GIACOMO MICHELI[‡]
(*Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK*)

and RETO SCHNYDER[§]
(*Institute of Mathematics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*)

[Received 14 August 2017]

Abstract

Let q be an odd prime power and D be the set of irreducible polynomials in $\mathbb{F}_q[x]$ which can be written as a composition of degree two polynomials. In this paper, we prove that D has a natural regular structure by showing that there exists a finite automaton having D as accepted language. Our method is constructive.

1. Introduction

It has been of great interest in recent years the study of irreducible polynomials which can be written as composition of degree two polynomials (see for example [1, 2, 6, 8, 9, 11, 13, 14, 15–17]). Such polynomials are also used in other contexts, see for example Rafe Jones' construction of irreducible polynomials reducible modulo every prime [15] or the proof of [3, Conjecture 1.2] in [5]. In this paper, we explain how the theory of this class of polynomials completely fits in a general context which allows the use of powerful machinery coming from the theory of finite automata (in characteristic different from 2). In fact, we show that some irreducibility questions over finite fields can be translated into automata theoretical ones (see Definition 3.6 and Theorem 3.8). As a side result, we also obtain that the set of irreducible polynomials which can be written as the composition of degree two polynomials is naturally endowed with a regular structure given by Theorem 3.11.

Let $\mathbb{F}_q$ be a finite field of odd characteristic and let $\mathcal{S} \subset \mathbb{F}_q[x]$ be a set of monic degree two polynomials. In this paper, we consider the set $\mathcal{S}$ to be an alphabet, and a word $f_1 \cdots f_k \in \mathcal{S}^*$

[†]E-mail: af612@cam.ac.uk

[‡]Corresponding author. E-mail: giacomo.micheli@maths.ox.ac.uk

[§]E-mail: reto.schnyder@math.uzh.ch

corresponds to the composition $f_1 \circ \cdots \circ f_k \in \mathbb{F}_q[x]$. The empty word naturally corresponds to x. Let $\mathcal{I} \subset \mathcal{S}^*$ be the language of words whose corresponding compositions are irreducible. Our goal is to show that $\mathcal{I}$ is a regular language by providing an automaton that accepts it. The entire theory lifts to local fields under the assumption that the set $\mathcal{S}$ is finite and none of its elements has discriminant in the maximal ideal of the local field.

2. Distinguished sets and freedom

We include in this section some elementary facts concerning the freedom of the monoid generated by a finite set of irreducible degree two polynomials. These results will be needed in Section 3. Each polynomial $f \in \mathcal{S}$ can be uniquely written as $f = (x - a_f)^2 - b_f$ for some $a_f, b_f \in \mathbb{F}_q$. We define $D_{\mathcal{S}} = \{b_f : f \in \mathcal{S}\}$ to be the *distinguished set* of $\mathcal{S}$.

We denote by $\mathcal{S}^*$ the set of words over the alphabet $\mathcal{S}$, so $\mathcal{S}^*$ is the free monoid generated by the symbols in $\mathcal{S}$. Let $C_{\mathcal{S}} \subseteq \mathbb{F}_q[x]$ be the compositional monoid generated by $\mathcal{S}$, that is the smallest subset of $\mathbb{F}_q[x]$ containing $\mathcal{S}$ and x which is closed by composition.

We will denote by π the natural surjective morphism of monoids $\mathcal{S}^* \longrightarrow C_{\mathcal{S}}$ which maps a word $f_1 f_2 \ldots f_k \in \mathcal{S}^*$ to the composition $f_1 \circ f_2 \circ \ldots \circ f_k \in \mathbb{F}_q[x]$. Notice that, for an element $f \in \mathcal{S}$, we will denote by $f^{(n)}$ the n-fold composition of f with itself.

For $b \in D_{\mathcal{S}}$, we define A_b as the subset of all a in $\mathbb{F}_q$ such that there exists $f \in \mathcal{S}$ with $f = (x - a)^2 - b$. For any of the A_b, we define the difference set

$$A_b - A_b = \{a - a' : a, a' \in A_b\}.$$

We can define a relation $\sim$ on $\mathcal{S}^*$ by setting $u \sim w$ if there exists $\ell \in \bigcup_{b \in D_{\mathcal{S}}} (A_b - A_b)$ for which $\pi(u) + \ell = \pi(w)$. This relation is symmetric and reflexive but not transitive, unless $\bigcup_{b \in D_{\mathcal{S}}} (A_b - A_b)$ is an additive subgroup of $\mathbb{F}_q$.

In this section, we provide a computable condition to establish whether $C_{\mathcal{S}}$ is a free monoid, which will be needed later on.

PROPOSITION 2.1 *Let u, v be words of $\mathcal{S}^*$ of equal length $n \geq 1$. Let u', v' be the $(n-1)$-suffixes of u and v, respectively. Then*

(i) $\pi(u) = \pi(v)$ *implies* $u' \sim v'$,
(ii) $u \sim v$ *if and only if* $\pi(fu) = \pi(gv)$ *for some* $f, g \in \mathcal{S}$.

Proof. Let us first prove (i). Suppose $\pi(u) = \pi(v)$ and let us write

$$\pi(u) = (h_1 - a)^2 - b = (h_2 - a')^2 - b' = \pi(v)$$

for $h_1 = \pi(u')$, $h_2 = \pi(v')$. Then we have $(h_1 - a - h_2 + a')(h_1 + h_2 - a - a') = b - b'$. Since $h_1 + h_2 - a - a'$ has positive degree, this forces $b = b'$ and $h_1 - a - h_2 + a' = 0$. Now it is clear that $a, a' \in A_b$, which implies $a' - a \in A_b - A_b$, and then $u' \sim v'$.

Let us now prove (ii). If $u \sim v$, by definition we have $\pi(u) - a = \pi(v) - a'$ for $a - a' \in A_b - A_b$ for some b. Now, by squaring and subtracting b on both sides of the equality we get $f(\pi(u)) = g(\pi(v))$ for some $f, g \in \mathcal{S}$, and hence $\pi(fu) = \pi(gv)$. Vice versa, if there exists $f, g \in \mathcal{S}$ such that $\pi(fu) = \pi(gv)$, then (i) applies. □

LEMMA 2.2 *Let u, v be words of $\mathcal{S}^*$ of equal length $n \geq 1$. If $|D_{\mathcal{S}}| = |\mathcal{S}|$ or $|D_{\mathcal{S}}| = 1$, then we have that $u \sim v$ if and only if $\pi(u) = \pi(v)$.*

Proof. One direction is trivial: if $\pi(u) = \pi(v)$, then $u \sim v$. For the other direction, we look at the two cases separately.

In the case $|D_{\mathcal{S}}| = |\mathcal{S}|$, it follows from $u \sim v$ that $\pi(u) - \pi(v) = c \in A_b - A_b$ for some $b \in D_{\mathcal{S}}$. However, A_b consists of only one element, so $c = 0$.

For the case $|D_{\mathcal{S}}| = 1$, assume that $u \sim v$, so $\pi(u) = \pi(v) + c$ for some $c \in \mathbb{F}_q$. Let u', v' be the $(n-1)$-suffixes of u and v. Then, since $|D_{\mathcal{S}}| = 1$, we have that $(\pi(u') - a_1)^2 - (\pi(v') - a_2)^2 = c$ for some $a_1, a_2 \in \mathbb{F}_q$. As c is constant, this forces $(\pi(u') - a_1) - (\pi(v') - a_2) = 0$, which in turn forces $c = 0$ and hence $\pi(u) = \pi(v)$. □

The following proposition shows that the freedom of the monoid is ensured whenever $D_{\mathcal{S}}$ is either maximal or minimal.

PROPOSITION 2.3 *If $|D_{\mathcal{S}}| = |\mathcal{S}|$ or $|D_{\mathcal{S}}| = 1$, then $C_{\mathcal{S}} \cong \mathcal{S}^*$.*

Proof. Clearly, a polynomial of degree two in $C_{\mathcal{S}}$ cannot have two distinct writings in terms of compositions. Let F be a polynomial in $C_{\mathcal{S}}$ of minimal degree with two different writings, that is, such that $F = \pi(fu) = \pi(gv)$ for $f, g \in \mathcal{S}$ and $u, v \in \mathcal{S}^*$ of positive length. From $\pi(fu) = \pi(gv)$, one deduces by Proposition 2.1 that $u \sim v$. Lemma 2.2 now gives $\pi(u) = \pi(v)$, which implies $u = v$ by the minimality of F. □

COROLLARY 2.4 *If $|\mathcal{S}| = 2$ then $C_{\mathcal{S}} \cong \mathcal{S}^*$.*

Proof. Immediate by observing that $|D_{\mathcal{S}}| = 1$ or $|D_{\mathcal{S}}| = |\mathcal{S}| = 2$. □

3. An automaton for irreducible compositions

3.1. *Capelli's Lemma*

In this subsection, we describe the basic tools needed to establish the main result. We start with a well-known result by Capelli, which gives a necessary and sufficient criterion to control the irreducibility of the composition of two polynomials.

LEMMA 3.1 (*CAPELLI'S LEMMA*) *Let K be a field and $f, g \in K[x]$ polynomials. Let $\beta \in \overline{K}$ be a root of g. Then, $g \circ f$ is irreducible over K if and only if g is irreducible over K and $f - \beta$ is irreducible over $K(\beta)$.*

We now use Capelli's Lemma to produce a simple ancillary result which will help us in what follows.

LEMMA 3.2 *Let $g \in \mathbb{F}_q[x]$ be monic and irreducible of even degree, and let $f = (x - a_f)^2 - b_f \in \mathbb{F}_q[x]$. Then, $g \circ f$ is irreducible if and only if $g(-b_f)$ is not a square in $\mathbb{F}_q$.*

Proof. Let $d = \deg(g)$, and let $\beta \in \mathbb{F}_{q^d}$ be a root of g. According to Lemma 3.1, $g \circ f$ is irreducible over $\mathbb{F}_q$ if and only if $f - \beta$ is irreducible over $\mathbb{F}_{q^d}$. Writing $f - \beta = (x - a_f)^2 - (b_f + \beta)$, this is equivalent to $b_f + \beta$ not being a square in $\mathbb{F}_{q^d}$. Let $N\colon \mathbb{F}_{q^d} \to \mathbb{F}_q$ be the norm map. If $\beta_1, \ldots, \beta_d$ are the roots of g, we have

$$N(b_f + \beta) = \prod_{i=1}^{d} (b_f + \beta_i) = (-1)^d \prod_{i=1}^{d} ((-b_f) - \beta_i) = g(-b_f),$$

since d is even. Now we can conclude, since $b_f + \beta$ is a non-square in $\mathbb{F}_{q^d}$ if and only if $N(b_f + \beta) = g(-b_f)$ is a non-square in $\mathbb{F}_q$. □

We are now ready to state one of the basic ingredients of the proof of the main theorem, which will allow us to 'push' irreducibility questions for compositions of degree two polynomials on a finite level.

PROPOSITION 3.3 *Let $f_1, \ldots, f_k \in \mathbb{F}_q[x]$ be monic polynomials of degree two. Write $f_i = (x - a_i)^2 - b_i$ for all i. Then, $f_1 \circ \cdots \circ f_k$ is irreducible if and only if all of the following are non-squares in $\mathbb{F}_q$:*

- b_1
- $f_1(-b_2)$

  ⋮
- $(f_1 \circ \cdots \circ f_{k-1})(-b_k)$.

Proof. Clearly, f_1 is irreducible if and only if b_1 is a non-square. The rest follows by inductive application of Lemma 3.2. □

3.2. *Brief interlude on Automata Theory*

In this subsection, we recall the basic results needed in the next section. For the definition of *deterministic finite automaton* (DFA) and *non-deterministic finite automaton* (NFA), we refer for example to [**12**, Chapter 2]. Since all the automata in the paper will have a finite set of states, we will often omit the word finite.

Let Σ be a set of symbols (an *alphabet*) and Σ^* be the set of words over Σ, that is the free monoid generated by it. Let us recall that a subset $\mathcal{L}$ of $\mathcal{S}^*$ is called a *language*. Let $\cdot$ be the usual binary operation in Σ^* (that is, concatenation of words) and * be the unary operation on languages defined by $\mathcal{L} \mapsto \mathcal{L}^*$ where $\mathcal{L}^*$ is the smallest submonoid of Σ^* containing $\mathcal{L}$ (in the context of languages, this operation is often called *Kleene star*).

A language is said to be *regular* if it is finite or can be expressed recursively starting from finite sets using the operations $\cup, \cdot, ^*$ (see [**12**] for more details). The following fact is well known.

THEOREM 3.4 *A language is regular if and only if it is accepted by a DFA.*

We will need the following fundamental results from the theory of Automata.

THEOREM 3.5 *If a language $\mathcal{L}$ is accepted by a DFA or an NFA, then it is regular.*

Roughly, the theorem above states that the accepted languages of NFAs do not generalize the notion of regular language.

We will also be using the notion of a *partial deterministic finite automaton*, which is the same as a DFA, except the transition function is actually a partial function. If, when reading a word, a required transition is not defined, the word is rejected. Clearly, a partial DFA is a special case of an NFA, so languages accepted by partial DFAs are also regular.

3.3. *Putting all together: building the automaton*

We first define a finite deterministic automaton $\mathcal{N} = \mathcal{N}(\mathcal{S})$ using the data contained in $\mathcal{S}$.

DEFINITION 3.6 The states of the automaton $\mathcal{N}(\mathcal{S})$ are given by the following:

- A special start state $\mathfrak{I}$. It is accepting.
- For each $a \in \mathbb{F}_q$, we have a distinguished state $[a]$. It is accepting if $-a$ is a non-square.
- For each $a \in \mathbb{F}_q$, we have a state $\{a\}$. It is accepting if a is a non-square.

The transitions are as follows:

- For each $f \in \mathcal{S}$, we have a transition $\mathfrak{I} \xrightarrow{f} [-b_f]$.
- For each $f \in \mathcal{S}$ and each $a \in \mathbb{F}_q$, we have a transition $[a] \xrightarrow{f} \{f(a)\}$.
- For each $f \in \mathcal{S}$ and each $a \in \mathbb{F}_q$, we have a transition $\{a\} \xrightarrow{f} \{f(a)\}$.

REMARK 3.7 The reason we distinguish between the states $\{a\}$ and $[a]$ is that they may be accepting at different times: $\{a\}$ accepts if a is non-square, $[a]$ if $-a$ is non-square. In the case that -1 is a square in $\mathbb{F}_q$, the two are equivalent and we can identify the two types of states.

THEOREM 3.8 *The language $\mathcal{I}$ of irreducible compositions is regular.*

Proof. Let $\mathcal{L}$ be the regular language over the alphabet $\mathcal{S}$ that is accepted by the automaton $\mathcal{N}$ reading from *right to left*. It is easy to see that a single letter f is in $\mathcal{L}$ if and only if b_f is non-square. Furthermore, a word $f_1 \ldots f_k$, $k \geq 2$, is in $\mathcal{L}$ if and only if $(f_1 \circ \ldots \circ f_{k-1})(-b_{f_k})$ is non-square. By Proposition 3.3, it follows that the word $f_1 \ldots f_k$ corresponds to an irreducible polynomial if and only if each prefix $f_1 \ldots f_l$, $l \leq k$, lies in $\mathcal{L}$. In other words, $\mathcal{I}$ is the language of all words whose every prefix is in $\mathcal{L}$.

We want to prove that $\mathcal{I}$ is regular. To do so, let us first construct a non-deterministic automaton $\mathcal{U}$ from $\mathcal{N}$ by reversing all the arrows of $\mathcal{N}$ and setting the start state of $\mathcal{N}$ as final state of $\mathcal{U}$ and the final states of $\mathcal{N}$ as start states of $\mathcal{U}$. Observe that the accepted language of $\mathcal{U}$ is again the language $\mathcal{L}$, this time reading from left to right as usual. Now we use subset construction (see for example [**12**, Section 2.3.5]) to generate a new deterministic automaton $\mathcal{V}$ having the same accepted language as $\mathcal{U}$. Finally, we remove all non-accepting states from $\mathcal{V}$ to obtain the partial automaton $\mathcal{M}$. It is easily seen that $\mathcal{M}$ accepts exactly those words whose every prefix is accepted by $\mathcal{V}$, which as explained above is $\mathcal{I}$. □

REMARK 3.9 Notice that the language accepted by the final automaton $\mathcal{M}$ is prefix closed.

We now provide an example to see the theorem in action.

EXAMPLE 3.10 As a simple example, consider the case $q = 5$ and $\mathcal{S} = \{f, g\}$ with $f = x^2 - 2$ and $g = (x-1)^2 - 3$.

We first construct the interim automaton $\mathcal{N}$ using the method described in Definition 3.6. Since $p \equiv 1 \bmod 4$, we can identify the nodes $[a]$ and $\{a\}$. Note that we have removed the node $\{0\}$ since it is not reachable from $\mathfrak{I}$. The result is seen in Fig. 1.

After performing the transformation described in the proof of Theorem 3.8 and cutting out all unreachable states, we end up with the simple partial automaton $\mathcal{M}$ in Fig. 2. This shows that the irreducible compositions of f and g are precisely those of the form $f^{(n)}$, $f^{(n)} \circ g$, $f^{(n)} \circ g^{(2)}$ and $f^{(n)} \circ g^{(2)} \circ f$ for $n \geq 0$.

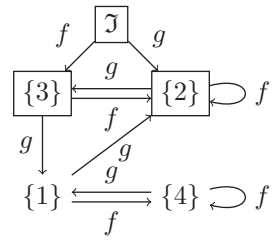

Figure 1. The interim automaton $\mathcal{N}$ for Example 3.10, coming from Definition 3.6. Boxed states are accepting.

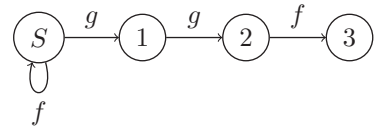

Figure 2. The automaton $\mathcal{M}$ accepting $\mathcal{I}$ for Example 3.10. All states are accepting.

Using the machinery we developed in the rest of the paper, we describe an infinite set of primes of $\mathbb{F}_q[x]$ having a finite regular structure.

THEOREM 3.11 *Let $\mathbb{F}_q$ be a finite field of characteristic different from 2. The set of monic irreducible polynomials having coefficients in $\mathbb{F}_q$ which can be written as a non-empty composition of degree 2 polynomials can be partitioned into a finite disjoint union $\bigsqcup_{a\in\mathbb{F}_q}\mathcal{L}_a$ in such a way that each $\mathcal{L}_a$ is in natural bijection with the words of a regular language $\mathcal{L}$, which is independent of a. In particular, the set of monic irreducible polynomials has a finite regular expression in terms of the elementary operations $\cup,\cdot,^*$.*

Proof. Let D be the set of monic irreducible polynomials in $\mathbb{F}_q[x]$ that can be written as non-empty composition of degree 2 polynomials. Let $\mathcal{S} = \{x^2 - b : b \in \mathbb{F}_q\}$. By Proposition 2.3, $C_{\mathcal{S}}$ is isomorphic to $\mathcal{S}^*$, so it is naturally embedded in $\mathbb{F}_q[x]$. Apply now Theorem 3.8 to obtain the regular language of irreducible polynomials $\mathcal{I}$ generated by $\mathcal{S}$, and let $\mathcal{L} = \mathcal{I}\setminus\{x\}$. Let $\psi_a\colon \mathcal{L} \longrightarrow \mathbb{F}_q[x]$ be the shift map defined by $f(x) \mapsto f(x+a)$. Let $\mathcal{L}_a = \psi_a(\mathcal{L})$. It is easy to observe that for any polynomial $f \in D$, there exists $a \in \mathbb{F}_q$ such that $f(x-a)$ can be written as an element of $C_{\mathcal{S}}$. This shows that

$$D = \bigcup_{a\in\mathbb{F}_q} \mathcal{L}_a.$$

It remains to show that $\mathcal{L}_a \cap \mathcal{L}_b = \varnothing$ if $a \neq b$, the final result will follow immediately. We argue by induction on the length of the words in $\mathcal{L}$ (that is, the degree of the polynomials). Let $a, b \in \mathbb{F}_q$ with $a \neq b$ such that there exist two words $v, w \in \mathcal{L}$ of minimal length ℓ such that $\psi_a(v) = \psi_b(w)$. If $\ell = 1$, this is clearly impossible, so let us assume $\ell > 1$. We can write $f(v'(x+a)) = g(w'(x+b))$ for some $f, g \in \mathcal{S}$ and v', w' suffixes of v and w, respectively. Therefore, for some $k, j \in \mathbb{F}_q$, we have

$$v'(x+a)^2 - w'(x+b)^2 = (v'(x+a) - w'(x+b))(v'(x+a) + w'(x+b)) = k - j.$$

Since v', w' are monic and the characteristic of $\mathbb{F}_q$ is different from 2, then the degree of the polynomial $(v'(x+a)+w'(x+b))$ is greater than or equal to 2. This forces both $k=j$ and $(v'(x+a)-w'(x+b))=0$, which contradicts the minimality of ℓ. □

EXAMPLE 3.12 For an example demonstrating Theorem 3.11, take $q=3$ and $\mathcal{S}=\{f,g,h\}$ with $f=x^2$, $g=x^2-1$, $h=x^2-2$. From Proposition 2.3, $C_{\mathcal{S}}$ is free and isomorphic to $\mathcal{S}^*$. Applying the construction, we get the automaton shown in Fig. 3. We see that the irreducible polynomials in $C_{\mathcal{S}}$ are exactly x, h, $h\circ g\circ f^{(n)}$ for $n\geq 0$, and $h^{(2)}\circ k$ for $k\in C_{\mathcal{S}}$ arbitrary (possibly the identity). Applying Theorem 3.11, it follows that the set of irreducible polynomials in $\mathbb{F}_3[x]$ that can be written as a non-empty composition of degree 2 polynomials is precisely

$$\bigcup_{a\in\mathbb{F}_3}\Big(\{h(x+a)\}\cup\{h\circ g\circ f^{(n)}(x+a):n\geq 0\}\cup\{h^{(2)}\circ k(x+a):k\in C_{\mathcal{S}}\}\Big).$$

Let us now describe two implications of our result in the case in which q is small compared with n. Recall that the *iterated logarithm* of a positive real number x, denoted by $\log^* x$, is the number of times the logarithm function must be iteratively applied before the result is less than or equal to 1.

COROLLARY 3.13 *For fixed q, we can list all monic irreducible polynomials of degree 2^n which are compositions of degree 2 polynomials with complexity $O(q^n 2^n n 8^{\log^*(2^n)})$, where the implied constant depends only on q.*

Proof. First, we choose $\mathcal{S}=\{x^2-a:a\in\mathbb{F}_q\}$ and write down the automaton $\mathcal{M}$ given by Theorem 3.8. The complexity of this step is $O(1)$, where the implied constant clearly depends only on q. Since the distinguished set is maximal, we can identify the polynomials $C_{\mathcal{S}}$ and the words in $\mathcal{S}^*$. Now the construction is recursive: let $\mathcal{L}_0(m)$ be the set of words of $\mathcal{S}^*$ of length m (so polynomials of degree 2^m). For any word $g\in\mathcal{L}_0(m)$ and any $f\in\mathcal{S}$, check if the word gf is accepted by the automaton $\mathcal{M}$: if it is, gf is an element in $\mathcal{L}_0(m+1)$. This check takes constant time for fixed q if for each element of $\mathcal{L}_0(m)$ we also store the state in which the automaton ends after reading it. As we observed in Remark 3.9, the language accepted by $\mathcal{M}$ is prefix closed, so all the elements of $\mathcal{L}_0(m+1)$ can be constructed in this way. Then, constructing $\mathcal{L}_0(m+1)$ from $\mathcal{L}_0(m)$ costs at most $O(q^{m+1})$.

It is elementary now to observe that $\mathcal{L}_0(1)$ can be easily constructed by selecting the irreducible degree 2 polynomials (again $O(1)$ operations) and, therefore, that $\mathcal{L}_0(n)$ can be constructed in time at most $O(q^n)$. Using the proof of Theorem 3.11 directly, we know that the set $D(n)$ of all

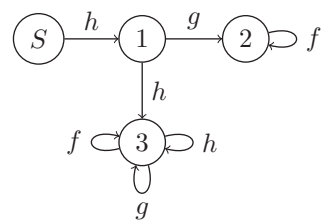

Figure 3. The automaton $\mathcal{M}$ accepting $\mathcal{I}$ for Example 3.12. All states are accepting.

irreducible polynomials of degree 2^n which are compositions of degree 2 polynomials can be written as

$$D(n) = \bigsqcup_{a \in \mathbb{F}_q} \mathcal{L}_a(n),$$

where $\mathcal{L}_a(n) = \{g(x+a) \colon g \in \mathcal{L}_0(n)\}$.

In order to represent the elements of $D(n)$ by coefficients, we now need to evaluate the elements of $\mathcal{L}_0(n)$. For a single such element, if we evaluate it from inside out, this means squaring polynomials of degree $d = 1, 2, \ldots, 2^{n-1}$ and subtracting a constant each time. According to [10], such a squaring can be done in time $O(d \log(d) 8^{\log^*(d)})$, the total time for each element hence being $O(2^n n 8^{\log^*(2^n)})$. □

REMARK 3.14 The reader should notice that this is much quicker than listing such polynomials by using an irreducibility test. This would have in fact complexity $O(q^n 2^n n 8^{\log^*(2^n)} I(2^n))$, where $I(2^n)$ is the cost of the chosen irreducibility test for a polynomial of degree 2^n. Again, the implied constant depends on the time of constructing the automaton, which just depends on the parameter q.

Another interesting fact is that (again for fixed q) we have an efficient deterministic algorithm to test irreducibility for polynomials which are a composition of degree two polynomials.

COROLLARY 3.15 *Let f be a polynomial of degree 2^n which is known to be a composition of monic degree two polynomials. Then f can be tested for irreducibility in time $O(2^n n^2 8^{\log^*(2^n)})$, where the implied constant depends only on q.*

Proof. We need to write f in the form $f = g_1 \circ g_2 \circ \ldots \circ g_n \circ \ell$, where $g_i = x^2 - a_i$ and $\ell = x - b$, with a_i and b in $\mathbb{F}_q$.

For this, suppose we have $F = G \circ H$ with $G = x^2 - a$, where $a \in \mathbb{F}_q$, and H is a polynomial of degree $d \geq 1$. Assume F is known, and we seek G and H. We apply the algorithm `Univariate decomposition` from [7, Section 2] with $r = 2$. This gives us the unique $\tilde{G}$ of degree 2 and $\tilde{H}$ of degree d with $\tilde{H}(0) = 0$ such that $F = \tilde{G} \circ \tilde{H}$. Clearly, setting $c = H(0)$, we have $\tilde{H} = H - c$ and $\tilde{G} = G(x + c) = x^2 + 2cx + c^2 - a$. From this, it is easy to recover c, a, G and H. The algorithm `Univariate decomposition` takes time $O(d \log(d)^2 8^{\log^*(d)})$, again using the multiplication algorithm from [10].

Applying this repeatedly to the f from the theorem will find the decomposition $f = g_1 \circ g_2 \circ \ldots \circ g_n \circ \ell$ in $O(2^n n^2 8^{\log^*(2^n)})$ time, which can then be checked for irreducibility with the automaton in time $O(n)$. Again, for fixed q, constructing the automaton has constant complexity, independently of the degree of the polynomial we are testing. □

REMARK 3.16 For fixed q, testing irreducibility for a polynomial of degree 2^n using for example Rabin's test [18] with fast polynomial operations costs $O(n4^n)$.

REMARK 3.17 As we already mentioned, the whole point is that our algorithm is very efficient in the regime of small fixed q and large n. Let us nonetheless have a quick look at the complexity with regard to q. If we follow the algorithm as described in Corollary 3.15 directly, the complexity appears to be exponential in q. In particular, we can expect the subset construction step to take exponential time and space.

Fortunately, it is not in fact necessary to construct the entire automaton to execute the above algorithm. Instead, we can solely construct the interim automaton $\mathcal{N}$ from Theorem 3.8. This takes

$O(q^2)$ field operations, and has to be done only once for each q. Then, given the decomposition of the polynomial into $f_1 \circ \ldots \circ f_n$ we can run a word $f_1 \ldots f_n$ through $\mathcal{N}$ directly as follows: define S_0 as the set of accepting states of $\mathcal{N}$. Then, for i from 1 to n, let S_i be the set of all states t such that there is an $s \in S_{i-1}$ and a transition $t \xrightarrow{f_i} s$ in $\mathcal{N}$. If at some point S_i does not contain the initial state $\mathfrak{I}$, we reject the word. Otherwise, we accept.

This method mirrors the reversal and subset construction from Theorem 3.8, except that only the parts that are actually used are computed. The complexity of this algorithm is easy to determine: when computing S_i from S_{i-1}, we can iterate over all states t of the automaton and check whether the unique outgoing transition labelled f_i ends in S_{i-1}. Hence, each step takes only $O(q)$ field operations, for a total cost of $O(q^2 + nq)$. Since for small q the quantity $2^n n^2 8^{\log^*(2^n)}$ dominates $q^2 + nq$, the complexity in terms of $\mathbb{F}_q$-operations remains unchanged.

4. Irreducible compositions over local fields

In this final section, we will show how the results of the previous sections can be lifted, under some additional hypothesis, to polynomials over local fields. Let K be a non-Archimedean local field with finite residue field $\mathbb{F}_q$ of odd characteristic. Let $\mathcal{O}_K$ be its ring of integers and ϖ be a uniformizer. We will denote by $\tilde{}$ the reduction map $\mathcal{O}_K[x] \to \mathbb{F}_q[x]$. Let us start by recalling the following lemma, which we state in a weaker form, sufficient for our purposes.

LEMMA 4.1 *Let L be any field and let $f, g \in L[x]$ be monic polynomials with $f = (x - a_f)^2 - b_f$ for some $a_f, b_f \in L$. Then we have*:

$$\operatorname{disc}(g \circ f) = \pm \operatorname{disc}(g)^2 \cdot 4^{\deg g} \cdot g(-b_f).$$

Proof. See [**14**, Lemma 2.6]. □

THEOREM 4.2 *Let $f_1, \ldots, f_k \in \mathcal{O}_K[x]$ be monic polynomials of degree 2 such that $\varpi \nmid \operatorname{disc}(f_1)$. Then $f_1 \circ \ldots \circ f_k$ is irreducible in $K[x]$ if and only if $\widetilde{f_1} \circ \ldots \circ \widetilde{f_k}$ is irreducible in $\mathbb{F}_q[x]$.*

Proof. One direction is obvious, so let us assume that $f_1 \circ \ldots \circ f_k$ is irreducible. For every $i = 1, \ldots, k$, let $f_i = (x - a_i)^2 - b_i$ for some $a_i, b_i \in \mathcal{O}_K$. By Proposition 3.3, we need to show that the following elements are not squares:

- $c_1 := \widetilde{b_1}$
- $c_2 := \widetilde{f_1}(\widetilde{-b_2})$

  ⋮
- $c_k := (\widetilde{f_1} \circ \ldots \circ \widetilde{f_{k-1}})(\widetilde{-b_k})$.

First, suppose that $c_t = 0$ for some $t \in \{1, \ldots, k\}$. This implies that $\widetilde{f_1}$ has a root, and since by hypothesis the discriminant of $\widetilde{f_1}$ is non-zero, by Hensel's lemma we can lift such a root to a root of f_1. But then $f_1 \circ \ldots \circ f_k$ is clearly reducible, which is a contradiction. Thus we can assume that $c_i \neq 0$ for all $i \in \{1, \ldots, k\}$. Now let $t \in \{1, \ldots, k\}$ be such that c_t is a non-zero square. By Proposition 3.3, this implies that $\widetilde{f_1} \circ \ldots \circ \widetilde{f_t}$ is reducible. On the other hand, applying Lemma 4.1 recursively and using the definition of the c_i's we get that

$$\operatorname{disc}(\widetilde{f_1} \circ \cdots \circ \widetilde{f_t}) = u \cdot \prod_{i=1}^{t} c_i^{2^{t-i}} \neq 0,$$

where u is an appropriate power of 2 (up to sign). This proves that $\varpi \nmid \operatorname{disc}(f_1 \circ \ldots \circ f_t)$ and since $f_1 \circ \ldots \circ f_t$ is irreducible by hypothesis, it defines an unramified extension of K. It follows that $\widetilde{f_1} \circ \ldots \circ \widetilde{f_t}$ is irreducible (see for example [**4**, Chapter 7]), giving a contradiction. □

It is clear that the hypothesis that $\varpi \nmid \operatorname{disc}(f_1)$ is necessary for the claim to hold, since for example $x^2 - \varpi$ is irreducible in $K[x]$, while its reduction is reducible in $\mathbb{F}_q[x]$.

Given a finite set $\mathcal{S} \subseteq \mathcal{O}_K[x]$ of monic polynomials of degree two with unitary discriminant, Theorem 4.2 shows that irreducible compositions of the elements of $\mathcal{S}$ correspond bijectively to irreducible compositions of the elements of $\widetilde{\mathcal{S}} \subseteq \mathbb{F}_q[x]$. Therefore, if we consider $\mathcal{S}$ as an alphabet and $\mathcal{I}$ is the language of irreducible compositions of the elements of $\mathcal{S}$, we deduce immediately the following corollary.

COROLLARY 4.3 *The language $\mathcal{I}$ is regular.*

Proof. It is enough to apply Theorem 3.8 to the language of irreducible compositions of the elements of $\widetilde{\mathcal{S}}$. □

The above corollary essentially states that the theory we developed in the rest of the paper lifts entirely to local fields, at least in the case in which the elements in $\mathcal{S}$ have unitary discriminant. It would be interesting to understand what happens when this condition is not satisfied.

5. Further research

One of the natural questions arising from the results in the present paper is whether Theorem 3.11 can be generalized to higher degree polynomials. In fact any lift of such results to polynomials of degree three or more would be of great interest, in particular because the necessary and sufficient criterion by Boston and Jones [**16**] (and the subsequent results on the subject such as [**1**, **2**, **6**, **8**, **11**]) only exists in degree two. In the context of local fields, another interesting issue arising from Theorem 4.2 of this paper is the following: how can one include singular polynomials in the generating set $\mathcal{S}$? In fact, the condition on the discriminant seems to be essential. Another question arising from these results is whether it is possible to explicitly compute the generating function of the language of irreducible compositions just in terms of the coefficient of the generating polynomials. It is indeed possible to compute such a function just in terms of the obtained automata, but this would drop most of the available information. In particular, it seems that the key ingredient to address this issue is to understand the structure of the finite submonoid of maps from $\mathbb{F}_q$ to $\mathbb{F}_q$ which can be written as composition of degree two polynomials. More generally, many of the questions and constructions which were related just to the compositions of a single polynomial, now seem to naturally arise in this more general context, were the rigidity of finite automata theory provides assistance.

Funding

The first author was supported by Swiss National Science Foundation Grant no. 168459. The second author is thankful to Swiss National Science Foundation Grant nos. 161757 and 171248.

Acknowledgements

The authors are thankful to the anonymous referee for his suggestions, which greatly improved the content and the readability of the paper. We are especially thankful for suggesting a reference which boiled down the complexity in Corollary 3.15 from $O(4^n)$ to $O(2^n n^2 8^{\log^*(2^n)})$.

References

1. O. Ahmadi, A note on stable quadratic polynomials over fields of characteristic two. 2009. https://arxiv.org/abs/0910.4556.
2. O. Ahmadi, F. Luca, A. Ostafe and I. E. Shparlinski, On stable quadratic polynomials, *Glasg. Math. J.* **54** (2012), 359–369.
3. J. C. Andrade, S. J. Miller, K. Pratt and M.-T. Trinh, Special sets of primes in function fields, *INTEGERS* **13** (2013), 2.
4. J. W. S. Cassels, Local fields*, volume 3 of* London Mathematical Society Student Texts, Cambridge University Press, Cambridge, 1986.
5. A. Ferraguti and G. Micheli, On the existence of infinite, non-trivial F-sets, *J. Number Theory* **168** (2016), 1–12.
6. A. Ferraguti, G. Micheli and R. Schnyder, On sets of irreducible polynomials closed by composition, (Eds. S. Duquesne and S. Petkova-Nikova), Arithmetic of finite fields*, volume 10064 of* Lecture Notes in Computer Science, Springer, Cham, 2016, 77–83.
7. J. von zur Gathen, Functional decomposition of polynomials: the tame case, *J. Symbolic Comput.* **9** (1990), 281–299.
8. D. Gomez and A. P. Nicolás, An estimate on the number of stable quadratic polynomials, *Finite Fields Appl.* **16** (2010), 401–405.
9. D. Gómez-Pérez, L. Mérai and I. E. Shparlinski, On the complexity of exact counting of dynamically irreducible polynomials. 2017. https://arxiv.org/abs/1706.04392.
10. D. Harvey, J. Van Der Hoeven and G. Lecerf, Faster polynomial multiplication over finite fields, *J. ACM (JACM)* **63** (2017), 52.
11. D. R. Heath-Brown and G. Micheli, Irreducible polynomials over finite fields produced by composition of quadratics. *Revista Matemática Iberoamericana (to appear)*, 2017. ArXiv preprint: https://arxiv.org/abs/1701.05031.
12. J. E. Hopcroft, R. Motwani and J. D. Ullman, *Introduction to Automata Theory, Languages, and Computation*, Addison Wesley, Boston, MA, 2001.
13. R. Jones, Iterated galois towers, their associated martingales, and the p-adic mandelbrot set, *Compos. Math.* **143** (2007), 1108–1126.
14. R. Jones, The density of prime divisors in the arithmetic dynamics of quadratic polynomials, *J. Lond. Math. Soc.* **78** (2008), 523–544.
15. R. Jones, An iterative construction of irreducible polynomials reducible modulo every prime, *J. Algebra* **369** (2012), 114–128.
16. R. Jones and N. Boston, Settled polynomials over finite fields, *Proc. Am. Math. Soc.* **140** (2012), 1849–1863.
17. A. Ostafe and I. E. Shparlinski, On the length of critical orbits of stable quadratic polynomials, *Proc. Am. Math. Soc.* **138** (2010), 2653–2656.
18. M. O. Rabin, Probabilistic algorithms in finite fields, *SIAM J. Comput.* **9** (1980), 273–280.